Not your ordinary CDN

Andrei Vasilenkov

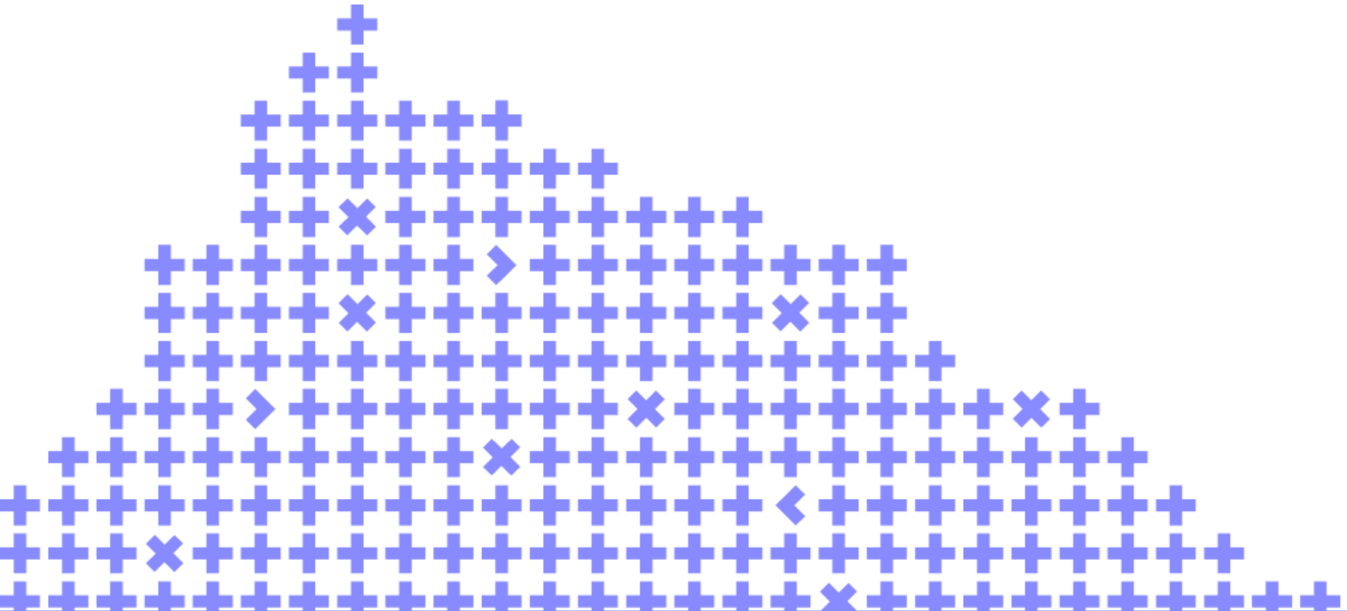

++ High Load Armenia

Co-organizer

Yandex

1. Streaming protocols
2. How to select CDN location?
3. Piecewise Linear Interpolation
4. PID Controller
5. Conclusion

Manifest / Playlist

Content type

Live event, VoD

Content duration

DRM settings

Content versions

Content versions

Video tracks

- Resolutions
- Bitrates
- Codecs

Audio tracks

- Language (dubbing, commentary track)
- Bitrates
- Codecs
- Number of channels

Subtitles / Captions

- Language
- Container

How does player interact with all of this?

- Filters content versions it is able to play
- Selects video/audio track by bitrate
- Chooses audio track version

  E.g. according to user’s language
- Builds URLs for media segments
- Downloads media segments and plays them seamlessly one by one

Considerations regarding CDN

- User session

  Media is played continuously, player repeatedly downloads new segments

- Response time matters

  Buffer can get empty if server fails to respond in time

  Buffer is limited when streaming Live content

- Heavy content

  Video segments’ size grows if you improve quality (usually)

Interaction with CDN

- Anycast

  strm.yandex.ru
- Anycast + HTTP 302 Redirect

  strm.yandex.ru → mskm901.strm.yandex.ru
- Dedicated CDN host

  mskm901.strm.yandex.ru

1. Streaming protocols
2. How to select CDN location?
3. Piecewise Linear Interpolation
4. PID Controller
5. Conclusion

Network is a large AS

- Tens of locations
- Hundreds of nodes
- Hundreds of network links

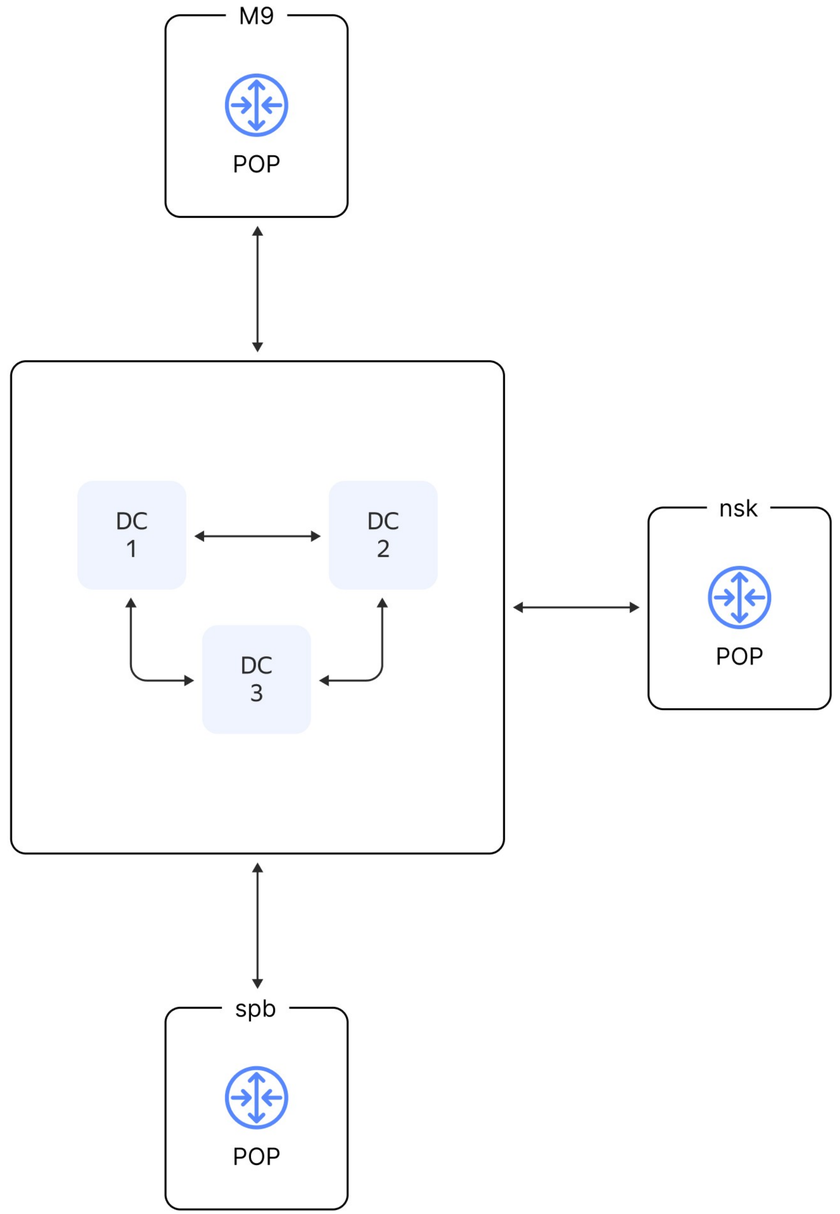

IP network prefix tree

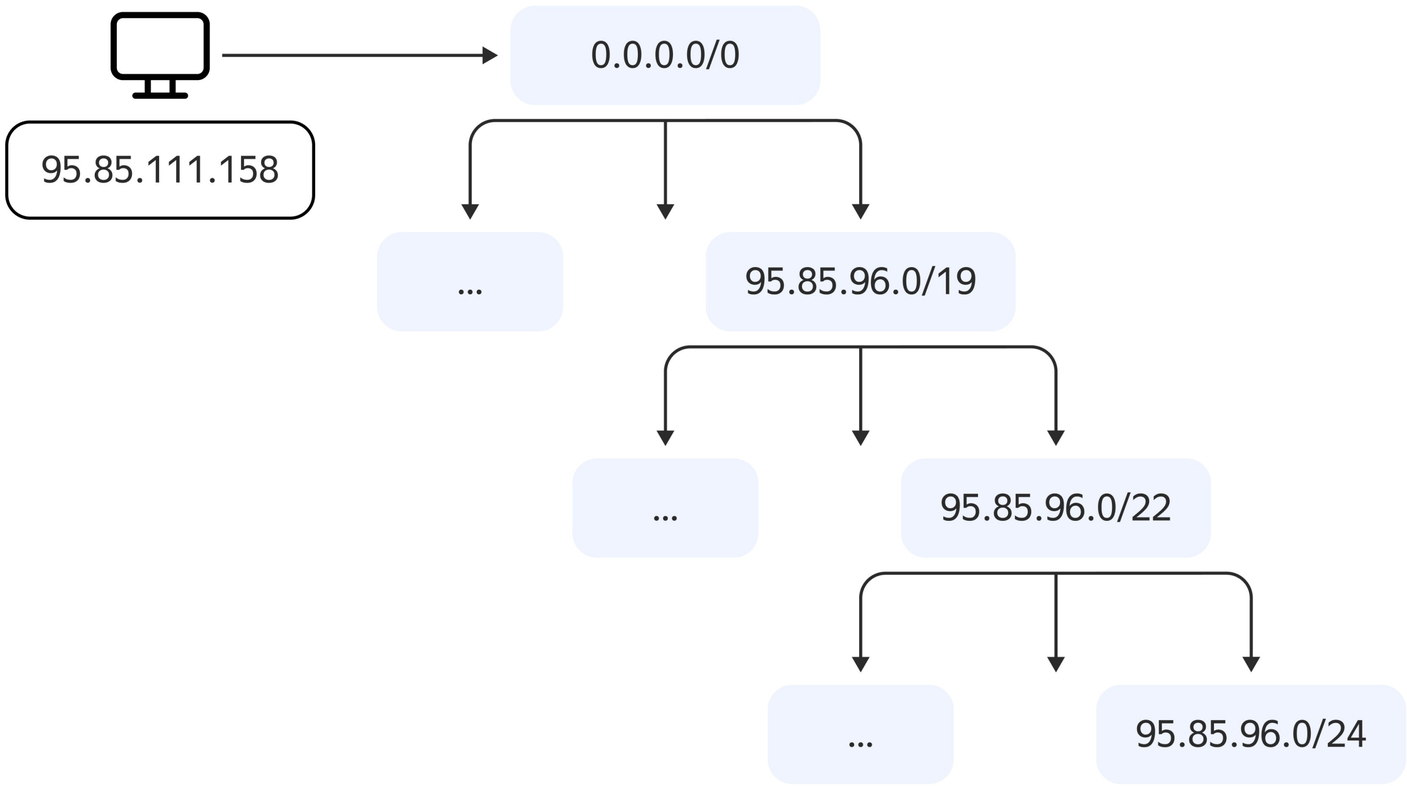

IP network prefix tree

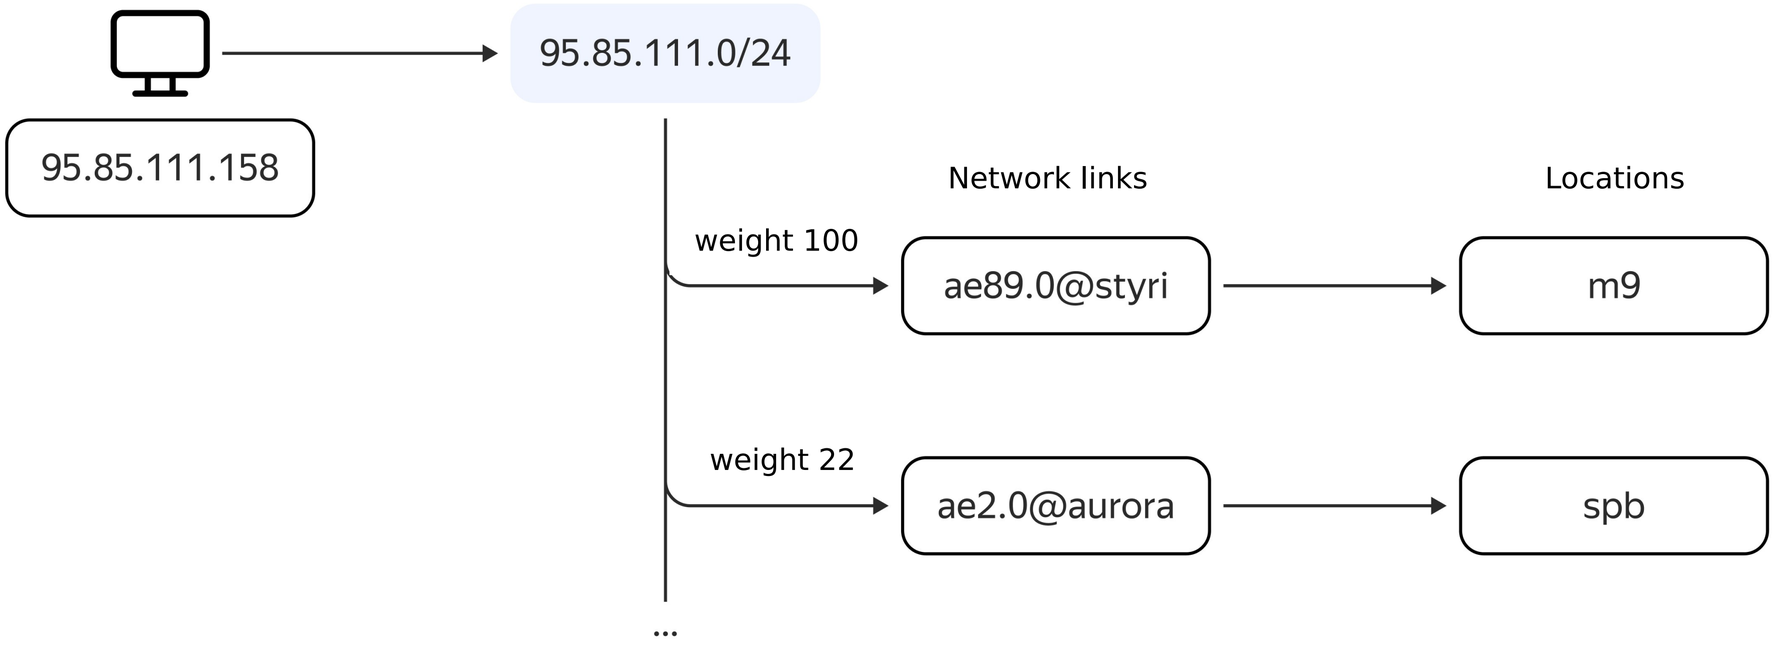

Network link weight

Weight helps to manage users' distribution based on

- Link bandwidth
- Network topology
- Server metrics

  E.g. RTT from tcp_info
- End user metrics

  Perf logs from players and browsers

Network link weight

The link is selected.
Now, can we use it?
Not right away!

- Weight is static for a long period
- Weight does not take current network load into consideration

Metrics required to use link

- Bandwidth
- Current status
  - Operational and available
  - Operational, but forbidden to be used
  - Out of service
- Current load
  - Gathered on network equipment — accurate but slow
  - Gathered on edge nodes — inaccurate but fast

How to use collected metrics

Balancing algorithm should consider

- Free bandwidth
- Link overload
- Link load waves
- Network utilization

User distribution algorithm

Link load management using two variables

- Slowdown — percentage of new user sessions that will not be serviced by chosen link
- Drop — percentage of current user sessions that should be moved away from chosen link

1. Streaming protocols
2. How to select CDN location?
3. Piecewise Linear Interpolation
4. PID Controller
5. Conclusion

Slowdown and drop

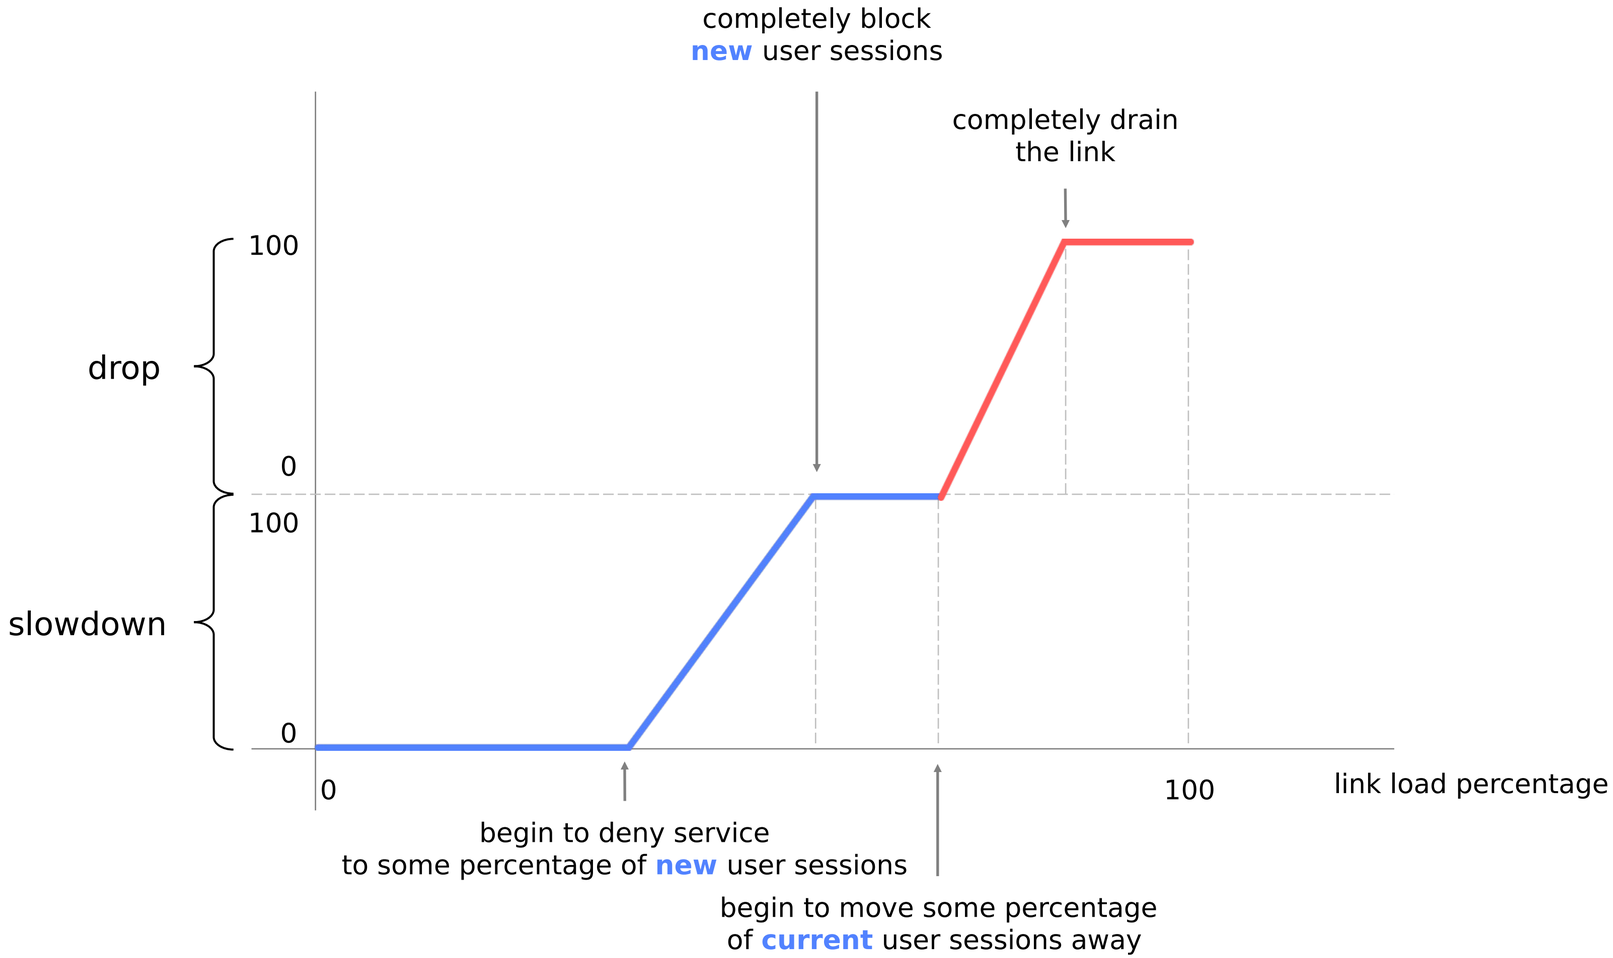

Load resonance

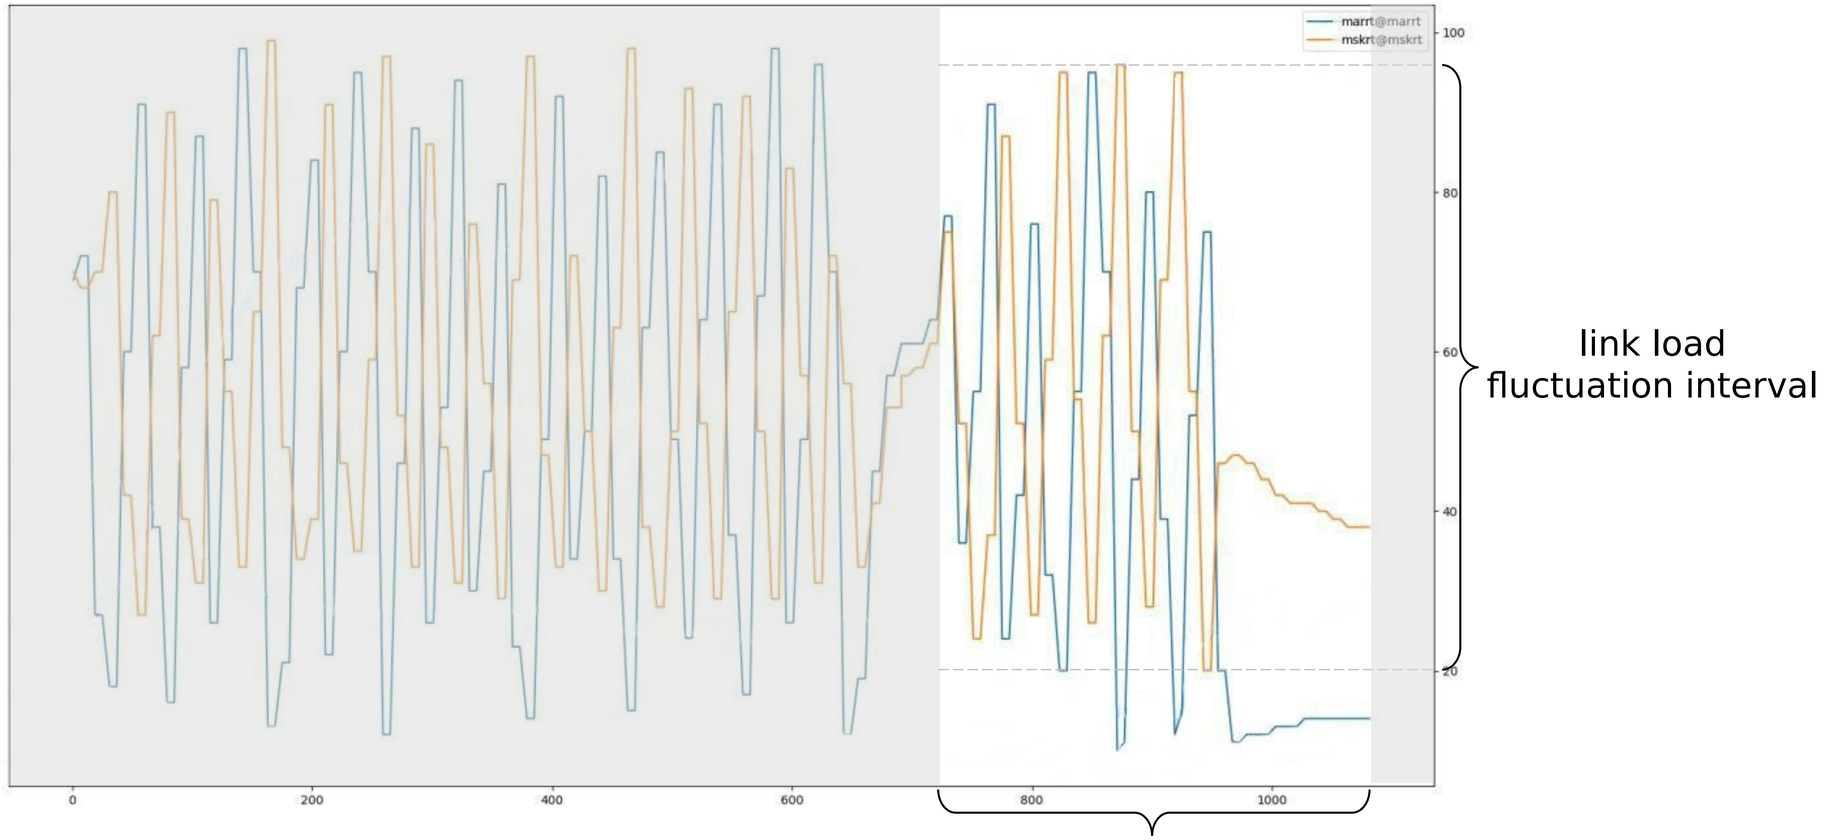

Taking history into account

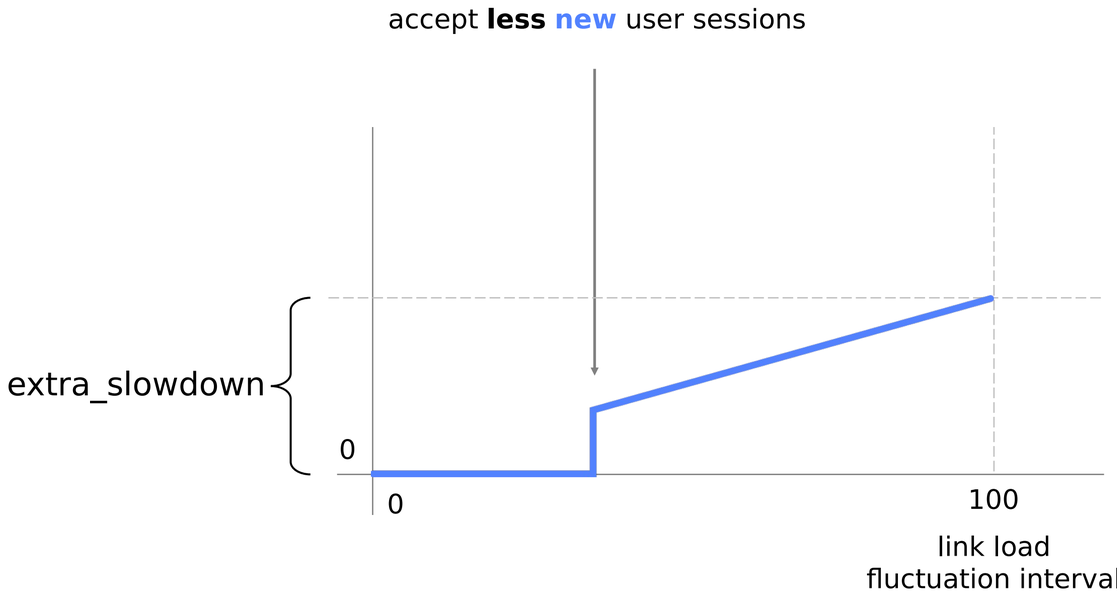

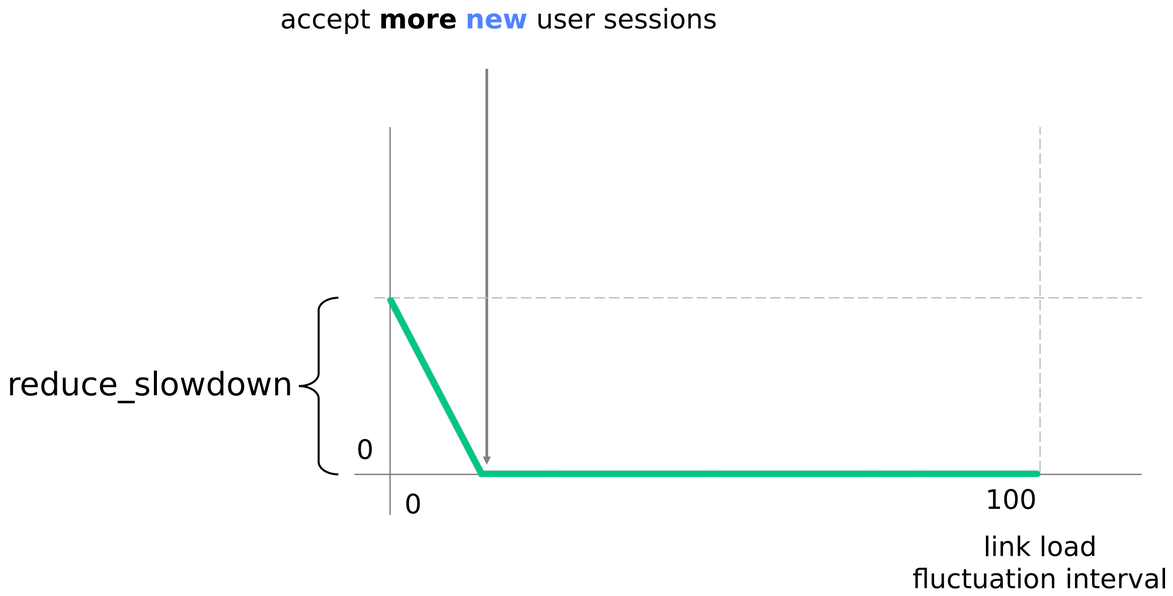

slowdown += extra_slowdown – reduce slowdown

Pros and cons

Takes available link bandwidth into account

Prevents link overload

Link load resonance and fluctuations

Can’t set desired link load target

Using collected metrics

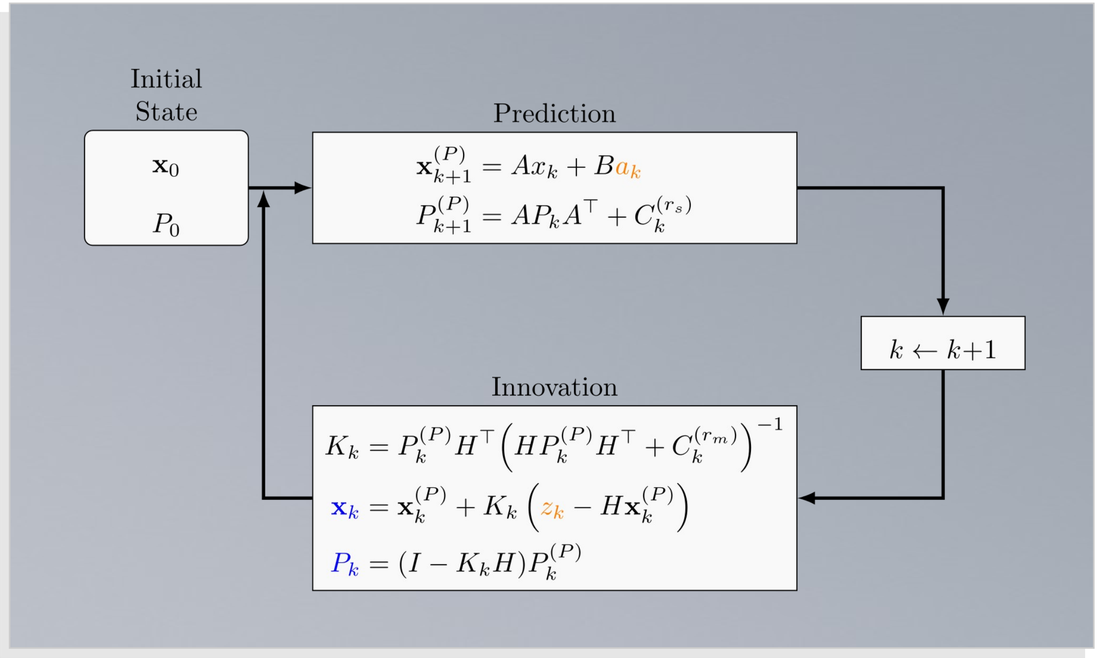

Kalman Filter

Makes predictions and observations, and then iteratively tries to correct them to remove the noise

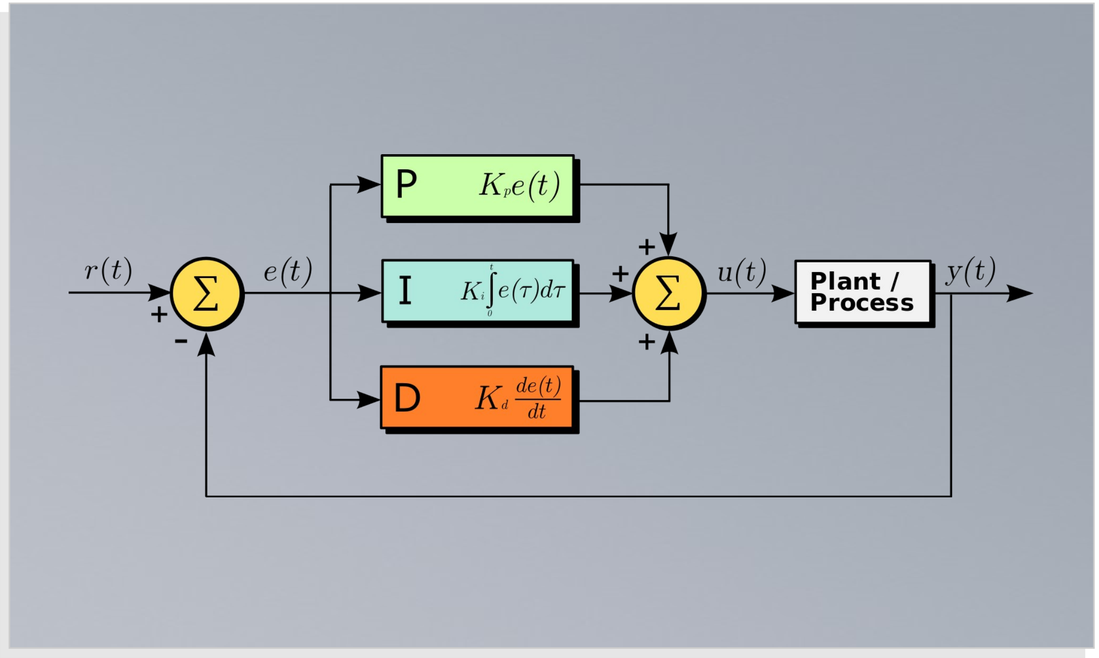

PID controller

Calculates error and tries to minimize it using current measurements and history data

1. Streaming protocols
2. How to select CDN location?
3. Piecewise Linear Interpolation
4. PID Controller
5. Conclusion

P Controller

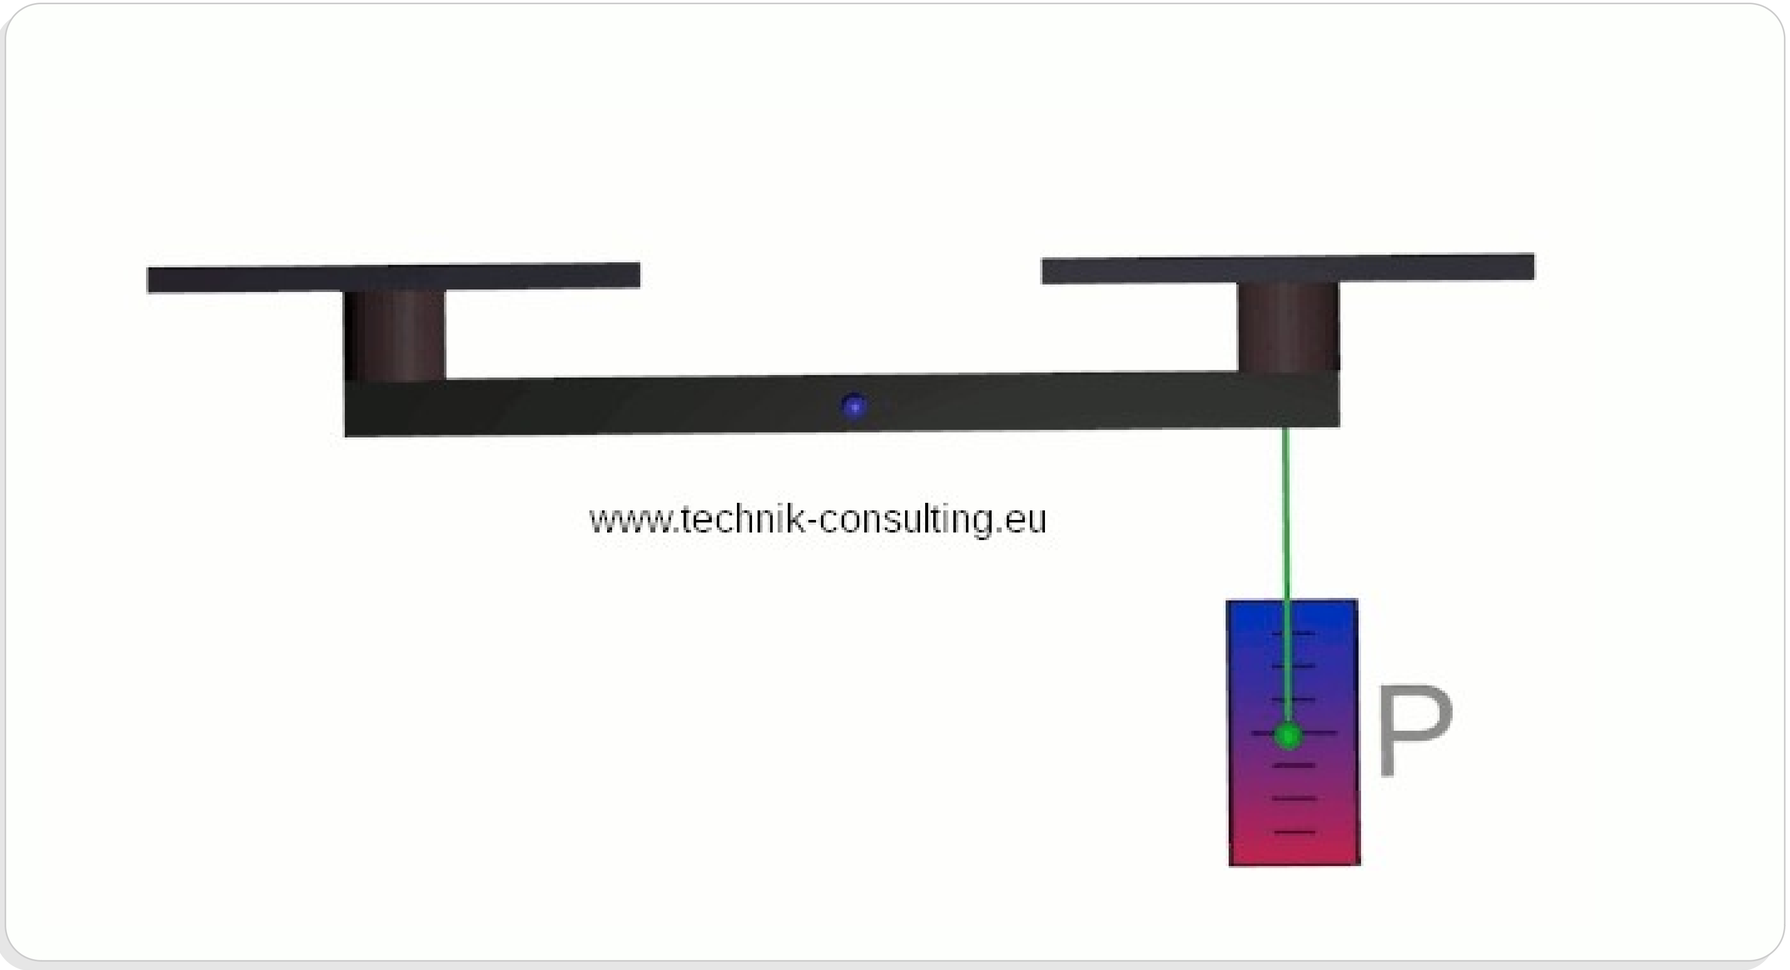

technik-consulting.eu/en/optimizing/drone_PID-optimizing.html

I Controller

technik-consulting.eu/en/optimizing/drone_PID-optimizing.html

I Controller

technik-consulting.eu/en/optimizing/drone_PID-optimizing.html

PID Controller

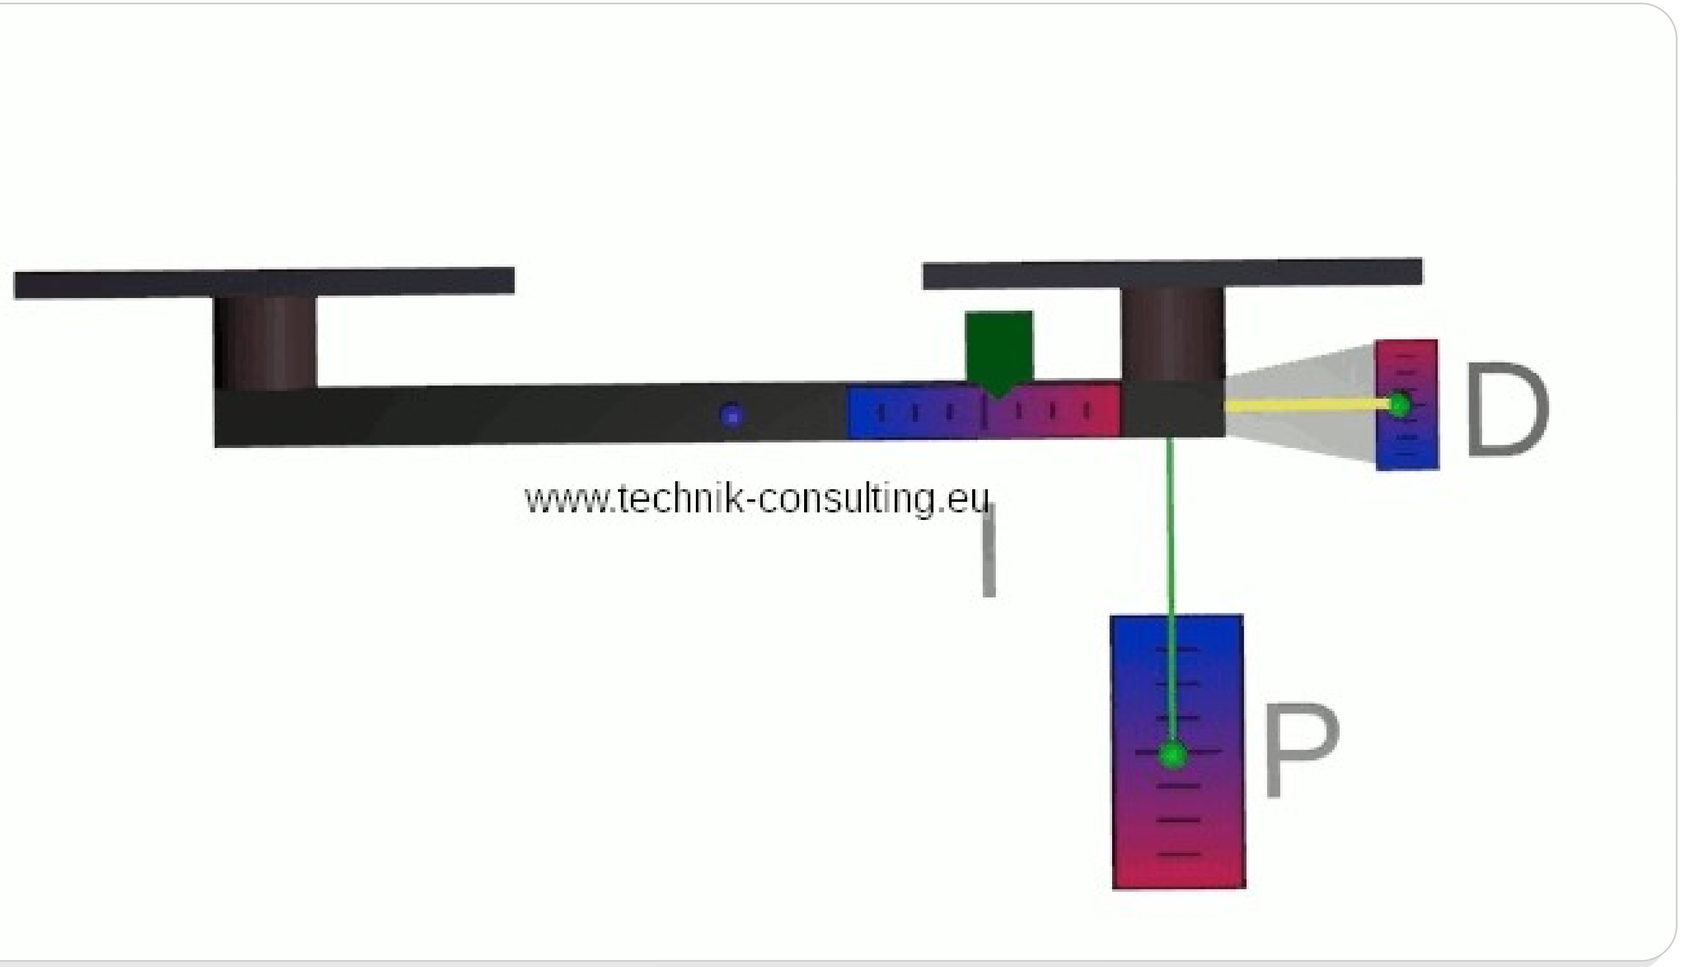

technik-consulting.eu/en/optimizing/drone_PID-optimizing.html

Requirements

Technical

- Fast signal delivery — collect link metrics as fast as possible
- Choose Kp, Ki and Kd values
- Select link load target value

Administrative

- Partial deploy of new algorithm
- Fast fallback
- Custom settings for certain links

And that’s how it’s done!

And that’s how it’s done!

PLI algorithm

- Works but not optimal
- Requires manual setup and tuning in real time

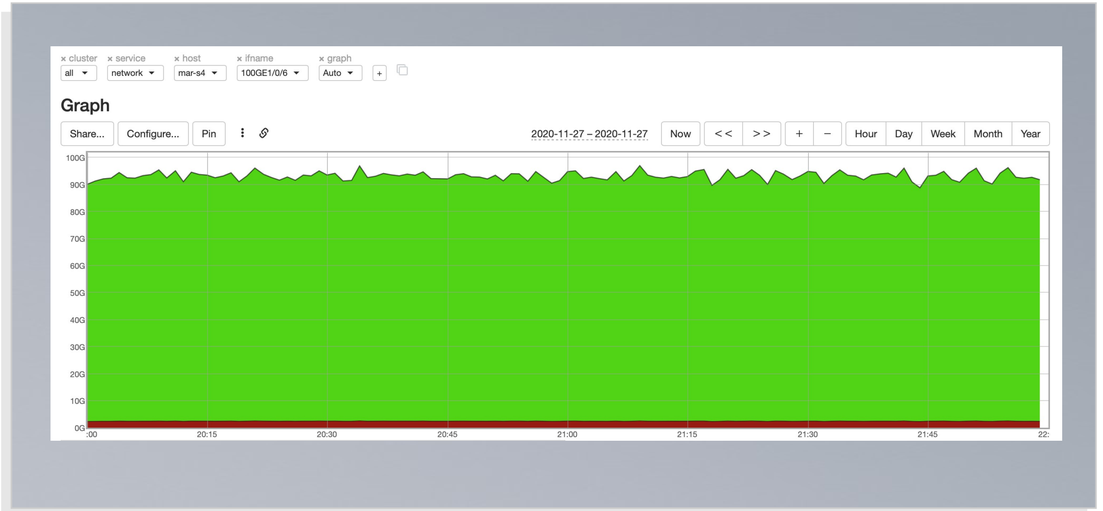

PID controller

- Works (almost) great
- Requires careful initial setup
- …but (almost) no supervision!

1. Streaming protocols
2. How to select CDN location?
3. Piecewise Linear Interpolation
4. PID Controller
5. Conclusion

CDN:
how to make it work

Balancing algorithm should consider

- Collect metrics
- Think about your case
- Know your theory
- Experiment
- Be ready to fallback
- Optimize your resources
- Have a great team!

Leave your feedback!

You can rate the talk and give a feedback on what you've liked or what could be improved

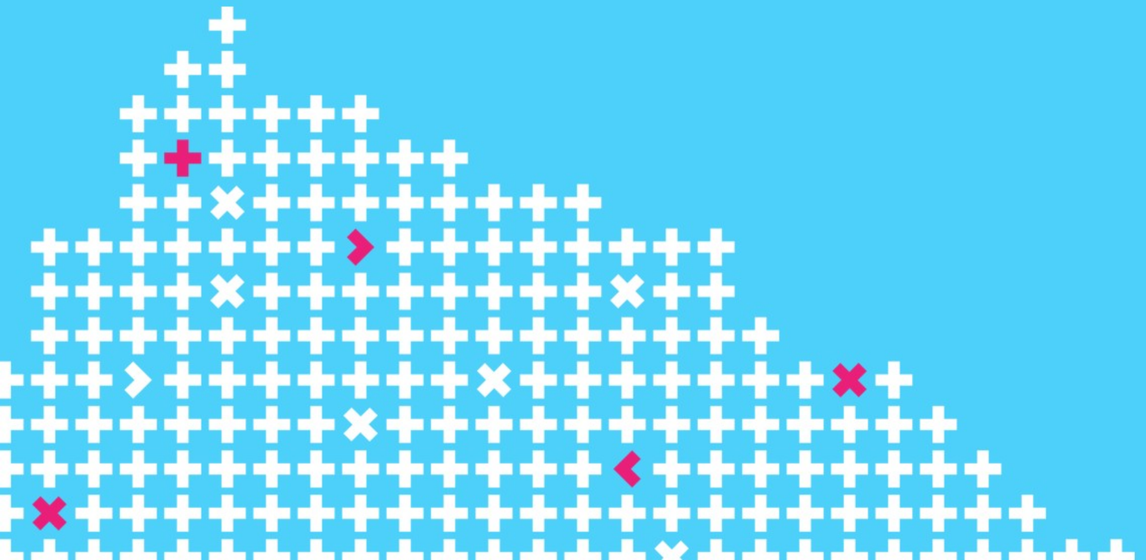

Co-organizer

Yandex